

SAM Cache Space Management Group Quotas

Igor Mandrichenko

Acknowledgements

- This proposal is the result of discussions with
 - Andrei Baranovsky
 - Gabriele Garzoglio
 - Igor Terekhov

Current Status, Problem Description

- Current cache space management mechanism is static hard quotas
- Inefficient resource utilization
 - Hard quotas may not “overlap” – “active” group can not preempt “passive”
 - Daniel Wicke: *With the current system however we have to manually reduce the cache size of 'dzero' to 70% and even manually remove 30% of the files from the cache, before we can assign 30% to 'd0production'*
 - When there is no room, cache removes files owned by the *same* group, never by *other* group

Proposal

- Basic idea: allow SAM Cache remove files owned by other groups – in *controlled, fair* way
- Introduce 2 quotas for each group:
 - Guaranteed quota
 - Relatively low, may not overlap
 - Maximum quota
 - May equal total capacity of the cache, overlaps are welcome
- Actual allocation dynamically balances between the two quotas

Fair Share

At any time, each *active* group has *target* allocation or fair share of space

$$T_i = C * F_i / \text{Sum}(F_j \text{ over all } \textit{active} \text{ groups})$$

C – total capacity of the cache

F_i – fair share factor – statically defined in the configuration

Algorithm

Any time there is no room for a file needed by a group:

1. Sort all (active and inactive) groups in by the difference between their current allocation (X) and their current target allocation (T):
$$D_i = X_i - T_i$$
2. Choose the group with greatest D which:
 - has at least one file which is not currently in use by any consumer and
 - with $X > \text{Guaranteed Quota}$
3. Use cache clean-up policy for the group to choose an unused file to be removed. Remove the file.
4. If more room is needed, go back to step 1

Details

- How to define *active* group ?
 - One definition: the one which has at least one Project Master running
 - Imperfect but seems to be good enough, at least to start with
- What about files shared by two or more groups ?
 - Currently ownership is transferred – good enough for the algorithm